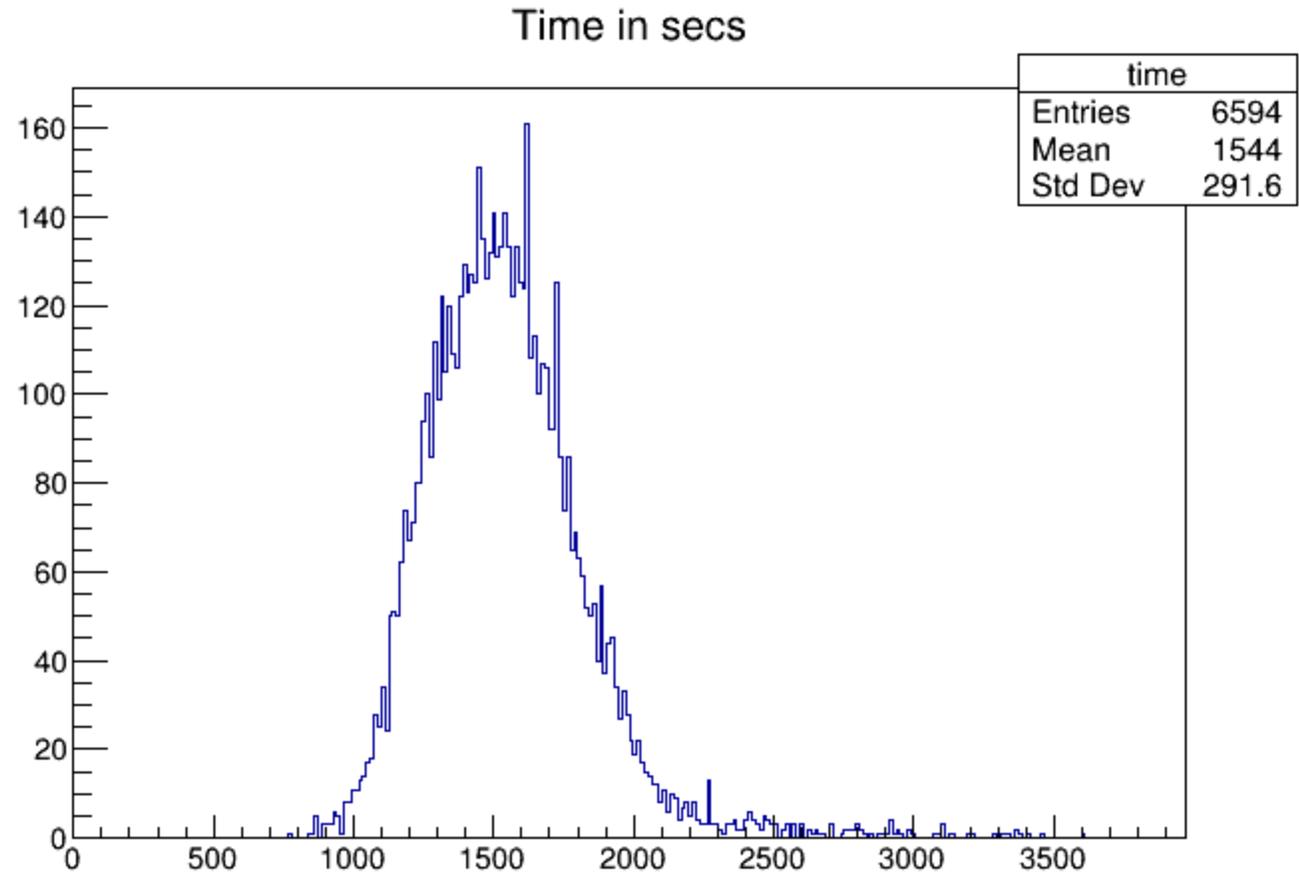


sPHENIX Status

Tracking

- Very active development
- Very good progress on Tracking speed
- More speedups going through Jenkins right now



Condor job length. Includes copying input and output from dCache and clustering (real data reco), 90 MinBias events with 50kHz pileup: 17sec/evt

MDC status

- Now completely on auto pilot – run by crons, you just kick off the initial G4 simulation, subsequent passes start when input ready
- 50kHz pileup added
- Not a single code related crash (one job died in G4)
 - Farm/condor issues kill jobs on low level (few per day)
- dCache can keep up – users do not see a difference analyzing files from gpfs or dCache
 - Use xrdcp to copy files but root's dccp interface (need to switch to xrootd)
- With dCache Users get trained to use the file catalog

MDC status

- 3 months of production running resulted in
- 1M MinBias AuAu events
- 1M semi central AuAu Events
- 1M central hijing events
- 100M HF pp events
- 6M MB pp events

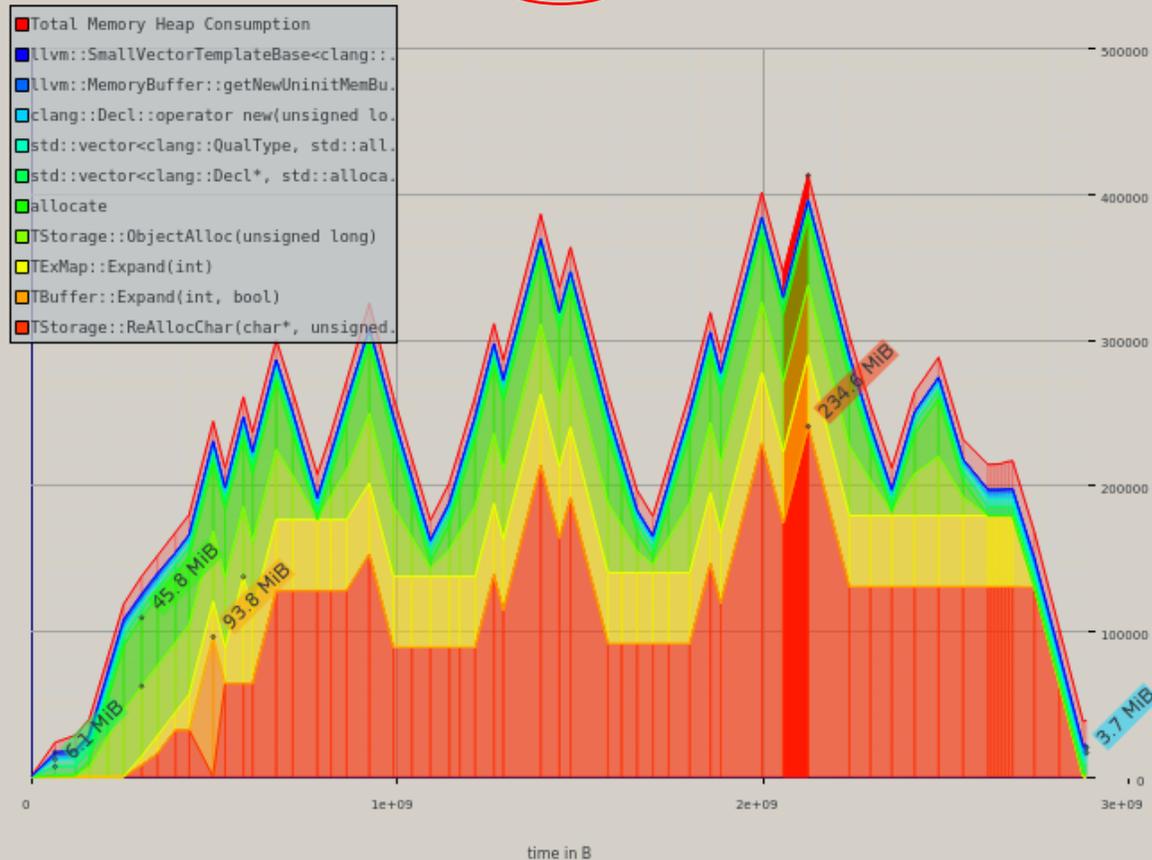
Memory for AuAu is a real problem, even pp needs 4GB.

Memory consumption – still suspect ROOT i/o

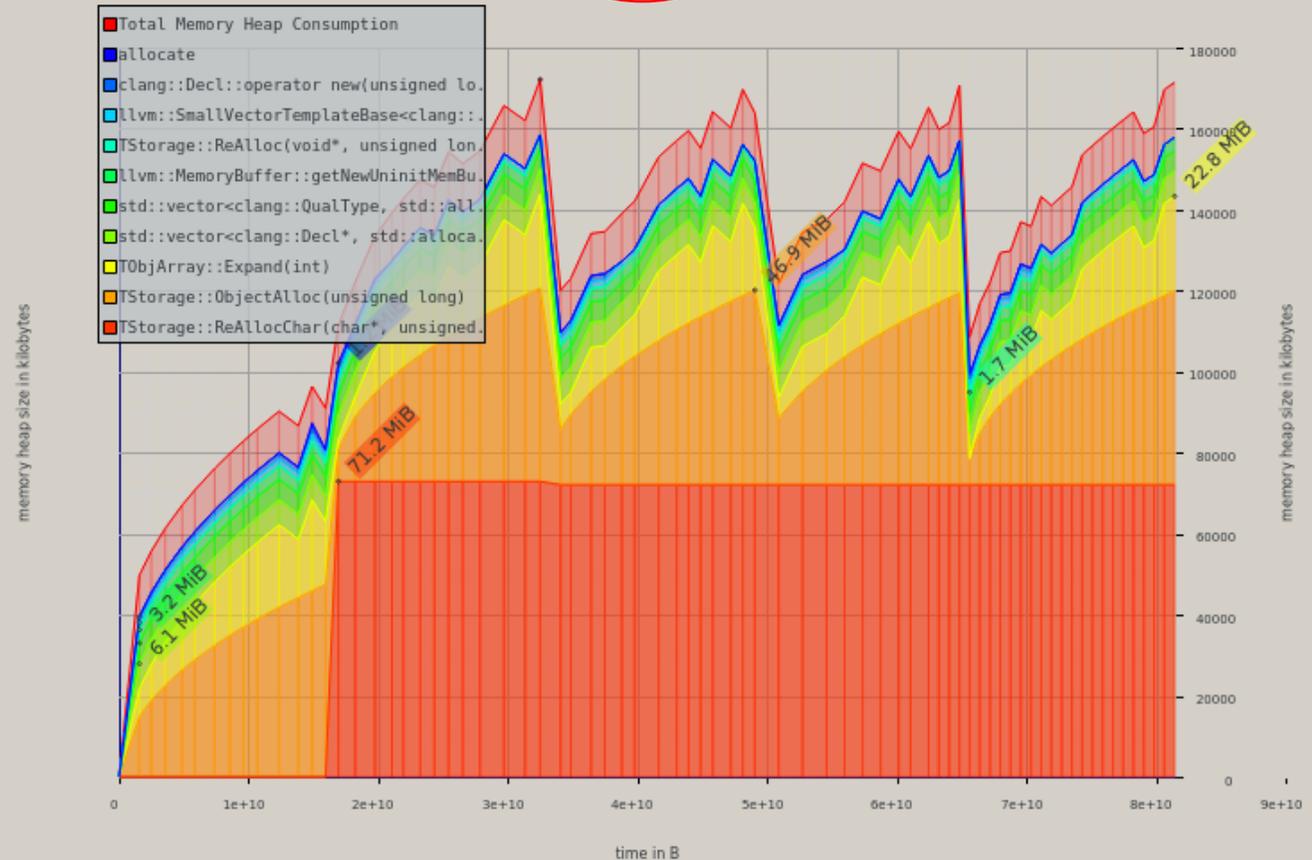
Test saving an stl map in a TBranch
(that's what we currently do)

Test saving a TClonesArray in a TBranch

Memory consumption of root.exe
Peak of 402.7 MiB at snapshot #41



Memory consumption of root.exe
Peak of 168.1 MiB at snapshot #26



Conditions DB

- Start to get clarity about the distortion corrections requirements
 - ALICE granularity 25ms – might go down to 10ms
- Before anyone panics
 - Our jobs should run about 24hours (not too long, not too short)
 - With 20sec/evt, that is $86400/20 \sim 4000$ events
 - Distributed over 2-3 passes – 8000-12000 events
 - 15kHz data rate \rightarrow jobs cover < 1 sec of data
 - Distortion correction extraction: 40-100 calibrations/job downloaded at end of job
 - Distortion corrected tracking: 40-100 calibrations/job – read at startup
 - 40-100 calibrations read (or written) in 24 hours

It is purely a question of scale, embarrassingly parallel, no complicated tags

We lost our DB expert, we can set up DBs but tuning or trouble shooting is beyond our current capabilities

PanDa

- Okay – I wrote a production system with DB backend over Christmas which will keep us afloat for the next mdcs
 - That buys us some time
- But no replacement for the real thing
 - PanDa need now driven by storage system testing
 - Needs PanDa server installation at BNL
 - dCache works nicely for analysis (automatic duplication of hot files, tape backend,...) – not necessarily the right choice for write once, read once then throw away

If we don't do this soon – the EIC will go their own way